# Lowering IrGL to CUDA

Sreepathi Pai Keshav Pingali
The University of Texas at Austin.
sreepai@ices.utexas.edu,pingali@cs.utexas.edu

arXiv:1607.05707v1 [cs.PL] 19 Jul 2016

## Abstract

The IrGL intermediate representation is an explicitly parallel representation for irregular programs that targets GPUs. In this report, we describe IrGL constructs, examples of their use and how IrGL is compiled to CUDA by the Galois GPU compiler.

***Keywords*** Irregular applications, amorphous data-parallelism, GPUs, compilers, CUDA

## 1. Compiling IrGL

Our implementation of the IrGL compiler is written in Python and operates on an AST of IrGL constructs (Listing 6). Apart from the constructs in Table 1, this AST also contains the CBlock construct for C++ code used in writing the operator. Our compiler parses this code for well-formedness as well as generating read/write sets, but is limited to C99 syntax and hence also accepts annotations to describe read/write sets. The compiler generates CUDA output, targeting Kepler and Maxwell GPUs. In the remainder of this report, we describe how the IrGL AST is lowered to CUDA. We assume a deep familiarity with CUDA. Listing 6 contains the definition of this AST in ASDL [10].

We use the following typographical conventions in this document – Terminals in the AST are indicated by a sans-serif font. Attributes on AST nodes, as well as values in code are represented by `typewriter` font.

## 2. Overall Structure of the AST

The IrGL AST is rooted at the Module node. The only supported children of Module are Kernel and global-level declarations. The Names node provides support for importing foreign names (such as `#define` constants) from C into the local IrGL symbol table. Many constructs in the AST, such as While, If, etc. only serve to expose control flow to the IrGL compiler. C code blocks, represented by CBlocks, must observe single-entry, single-exit control-flow behaviour, akin to basic blocks, but are allowed to call functions.

Not shown in the AST definition are compiler-specific *annotations* that can be applied to certain nodes. For example, the CUDA `__launch_bounds__` annotation conventionally used to indicate register-usage restrictions to the CUDA compiler is also supported by our IrGL compiler (which passes it through, but see Section 4.3). Other significant annotations will be mentioned when the nodes they annotate are discussed below.

## 3. Compiling Kernels

A Kernel node designates a *plain* IrGL kernel, a *host* kernel (`host=true`) or a *device* kernel (`device=true`). Host kernels execute on the CPU. Device kernels correspond directly to CUDA `__device__` kernels.

Only a host Kernel may use the IrGL orchestration constructs – Invoke, Iterate and Pipe. Similarly, only a non-host, non-device Kernel can use the IrGL kernel constructs – Atomic, Exclusive, Retry, Respawn, ReduceAndReturn and ForAll. Host kernels may use ForAll, but it is treated as For by our current compiler.

The compiler primarily uses device kernels when implementing the iteration outlining optimization. While user-provided device kernels are supported, they are treated as opaque and are largely ignored by our compiler.

The following subsections discuss how we compile the ForAll, Atomic, Exclusive and ReduceAndReturn kernel constructs. We defer discussion of Retry and Respawn to Section 4.

### 3.1 Compiling ForAll

The iterations of the outermost ForAll in an IrGL kernel are mapped to CUDA threads. Each CUDA thread usually executes multiple iterations. Section 4.3 describes the process the compiler uses to determine the number of CUDA threads to use for a kernel.

The object stored in `iterator` represents a random-access iterator and the order of iteration execution is not defined. By default, consecutive iterations of the ForAll are mapped to consecutive CUDA threads. However, other map-

[Copyright notice will appear here once 'preprint' option is removed.]

| Construct | Semantics |
|---|---|
| **Kernel Constructs** | |
| ForAll (*iterator*) { *stmts* } | Traverse *iterator* in parallel executing *stmts* |
| ReduceAndReturn (*bool-expr*) | Reduce values of *bool-expr* and return as kernel return value. The actual reduction, one of Any or All, is specified at kernel invocation. |
| Atomic (*lock-expr*) { *locked-stmts* } [ Else { *failed-stmts* } ] | Acquire *lock-expr* and execute *locked-stmts*. If an Else block provided, execute *failed-stmts* if *lock-expr* was not acquired. Implements divergence-free blocking locks [7]. |
| Exclusive (*object, elements*) { *locked-stmts* } [ Else { *failed-stmts* } ] | Try once to acquire locks for *elements* in *object* and execute *locked-stmts* on succeeding. On failure execute *failed-stmts* if provided otherwise execute next statement. One thread is guaranteed to execute *locked-stmts* on conflicts. |
| SyncRunningThreads | Compiler-supported safe implementation of GPU-wide global barriers [11] |
| Retry *item* (*or* Respawn *item*) | Push *item* into a retry worklist and re-execute the kernel. Use of Retry indicates a runtime conflict and triggers conflict management in the runtime (e.g. serial execution). |
| **Orchestration Constructs** | |
| [Any \| All () Invoke *kernel*(*args*) [)] | Invoke *kernel*, passing current worklists if kernel uses them. |
| Iterate [While \| Until Any \| All] *kernel*(*args*) [Initial (*init-iter-expr*)] | Iteratively invoke *kernel* until termination condition is met or worklist is depleted. if invoked standalone, establish fresh worklists using *init-iter-expr* for initialization, else pass current worklists. |
| Pipe [Once] { *stmts* } | Establish worklists to be used by *stmts*. Without Once, repeat *stmts* until worklists are empty. Nested Pipes will not establish worklists. |

**Table 1.** Summary of IrGL Statements, [] indicate optional parts, | indicates options. See Listing 6 for a formal definition of the abstract syntax tree.

pings are possible and these mappings are represented by compiler-specific annotations. For example, when implementing the Retry Backoff optimization, a mapping that distributes contiguous blocks of iterations to a single thread is used to reduce conflicts.

For is used to represent a loop whose iterations cannot be executed in parallel. When the amount of parallelism can only be discovered at runtime, as in most irregular graph algorithms, ForAll is used with additional synchronization inside the body of the loop.

### 3.2 Compiling Atomic and Exclusive

IrGL provides two statements – Atomic and Exclusive – that allow iterations of a ForAll loop to implement mutual exclusion. Both these statements implement functionality that are hard to get right [1, 7]. Atomic and Exclusive currently use software implementations but can be recompiled to use proposed hardware primitives [2, 7] if such primitives become available.

#### 3.2.1 Atomic

Atomic implements an atomic section, a block of code that is executed under control of a single lock. Atomic can be nested. Two forms of Atomic are supported, a default blocking form that waits for the lock to be acquired. The other form, indicated by a non-empty `fail_stmts` (i.e. Else), is non-blocking and executes the statements in `fail_stmts` when the lock was not acquired.

Atomic provides as a safe alternative to spinlocks, since spinlocks can deadlock on GPUs due to warp divergence. Internally, Atomic uses `atomicCAS` to set the state of a lock variable. If the `atomicCAS` fails to acquire the lock and the Atomic is blocking, a divergence-safe loop similar to that described in [5, 7] is generated by the compiler to reattempt locking.

We illustrate the use of Atomic using Borůvka's algorithm for minimum spanning tree. Borůvka's algorithm begins by treating each node of the input graph as a *component*. Then, it finds the minimum cross-component edge out of each components. These edges are added to the minimum spanning tree, and the components they connect are merged (or *unified*). The procedure then repeats on these merged components until only one component remains or no cross-component edges can be found (i.e. in a disconnected graph).

Particularly challenging is the implementation of finding the minimum edge out of a component. Since a component can consist of many nodes, recording the minimum edge requires at least two updates that must be carried out atomically – the minimum weight and the edge itself. No CUDA primitive suffices to perform multiple updates atomically. Previous implementations, notably that of Vineet et al. [9], store the weight and edge identifier as bitfields of a 32-bit integer, which would allow use of a single `atomicMin` at the cost of severely limiting the generality of the resulting code and its applicability to input graphs.

Listing 1 describes the `find-min-edge` kernel in the IrGL implementation that uses an Atomic to update the component's data in an atomic context. Note that this instance of Atomic is a blocking lock (i.e. no Else clause), so it can use

```
1  ForAll(nidx In wl) {
2    n = wl.pop(nidx);
3    n_component = components[n];
4    minwt = INF;
5
6    for(e In edges) {
7      // find minimum cross-component edge
8      // out of node n; store weight in minwt,
9      // and edge id in minedge
10   }
11
12   Atomic(component_locks[n_component]) {
13     if(component_minwt[n_component] > minwt) {
14       component_minwt[n_component] = minwt
15       component_minedge[n_component] = minedge
16       // other updates
17     }
18   }
19
20   if(node has cross-component edge) {
21     wl.push(n)
22   }
23 }
```

**Listing 1.** Find-Minimum-Cross-Component-Edge kernel of Borůvka's MST algorithm. The worklist initially contains all nodes.

a ticket lock from our runtime by simply setting a compiler flag.

### 3.3 Exclusive

Exclusive encloses a block of code that must acquire a large number of locks. Internally, threads are assigned priorities so that at least one thread always acquires all the locks it needs. Exclusive never blocks and may not be nested. The statements in `fail_stmt` are executed if the locks were not acquired.

The set of locks to be acquired is obtained from an array of lock indices. In the simplest form, `objs_lock` indicates an Array which contains the lock indices to be acquired. In the ArrayIterator form, `objs_lock` specifies an array iterator that yields the indices to be locked.

Consider the use of Exclusive in Delaunay Mesh Refinement (DMR). For DMR, the key kernel is `refine` which, when presented with a worklist of bad triangles, fixes each one of them in parallel. Each thread must have exclusive access to the triangles in the cavity of its bad triangle, as well as to the triangles that form the boundary of the cavity. The Exclusive construct is key to simplifying the implementation of DMR. Listing 2 illustrates how the triangles in the cavity are passed as input to Exclusive which then permits access to triangles in the cavity to one thread.

The Exclusive statement is implemented using a three-phase algorithm, with each phase separated by SyncRunningThreads. Our implementation is similar to the *race–prioritycheck–check* scheme described in [4]. In the first phase, Exclusive claims the locks supplied for the executing thread. If multiple threads claim the same lock, only one

```
1  ForAll(btidx In wl) {
2    bad_triangle = wl.pop(btidx);
3
4    build_cavity(bad_triangle, &cavity_size, &cavity);
5
6    Exclusive(mesh, cavity_size, cavity) {
7      delete_cavity(...)
8      // create new triangles
9    }
10
11   SyncRunningThreads();
12 }
```

**Listing 2.** Simplified Refine kernel in Delaunay Mesh Refinement.

claim is allowed. In the second phase, each thread checks to see if its claim for every lock stands. If another thread was granted the claim, then the threads that lost the claim attempt to win priority over the claim. In the third phase, each thread checks to see if it still retains the claims it sought. Any threads that do so proceed to execute the statements within the Exclusive, while those that do not move to the next statement or execute the Else clause if one is supplied. Our use of SyncRunningThreads for implementing Exclusive places restrictions on its use. An Exclusive must be placed in a location that will be uniformly executed by all threads and thus cannot be nested. In practice, this means that Exclusive may only be directly placed underneath the outermost ForAll.

### 3.4 Compiling SyncRunningThreads

Like CPUs, GPUs can create and execute many more threads than can run concurrently on hardware. However, unlike CPUs, a GPU thread usually runs to completion and cannot be preempted. Thus, the notion of a global barrier that synchronizes all threads does not readily translate to the GPU. If all threads being synchronized are not running concurrently, the global barrier will deadlock.

Nevertheless, barrier-like functionality is useful, even if it is limited to only those threads that are running concurrently. Such "device-wide barriers", have been described previously [11] and are supported in IrGL through the SyncRunningThreads statement, though our implementation is derived from code used in [3]. Safe use of SyncRunningThreads requires that a GPU kernel never be launched with more physical threads than can run concurrently. This number can vary from GPU to GPU and also depends on the size of the CUDA thread block.

CUDA 6.5 introduces the occupancy API that allows this number to be calculated at runtime for a kernel for each GPU present in the system. When our compiler generates code to launch an IrGL kernel that uses SyncRunningThreads, it limits the number of threads using this occupancy API to ensure deadlock-free execution. We note this method is not portable to other devices [8] and even on NVIDIA GPUs, assumes that all thread blocks of the kernel will eventually execute concurrently.

```
1  Kernel BFS(graph, LEVEL) {
2    ForAll(wlidx In wl) {
3      n = wl.pop(wlidx)
4      ForAll(e In graph.edges(n)) {
5        if(e.dst.level == INF) {
6          e.dst.level = LEVEL;
7          wl.push(e.dst.id);
8        }
9      }
10 }
11
12 LEVEL=0
13 Iterate BFS(graph, LEVEL) Initial [src] {
14   LEVEL++;
15 };
```

**Listing 3.** Level-by-level BFS kernel using a worklist

### 3.5 Compiling ReduceAndReturn

The ReduceAndReturn statement is used to construct a return value for an IrGL kernel using a reduction. The actual reduction is specified when invoking the kernel using Invoke or Iterate. Our compiler currently supports the Any and All reductions, and therefore the value to be reduced is a boolean expression stored in `value`.

Any returns `true` if any `value` evaluated to `true`. All returns `true` only if all `value` evaluated to `true`.

ReduceAndReturn terminates execution of the kernel, *except* when invoked inside a ForAll when it only terminates the current iteration.

The simplest compilation of ReduceAndReturn uses global memory storage and CUDA atomic instructions to implement these reductions. However, this can be made cheaper by re-using the cooperative conversion optimization machinery. In effect, each CUDA thread partially aggregates the ReduceAndReturn values, with atomics being used only to aggregate the values of each CUDA thread block. Unfortunately, since CUDA does not support virtual functions, we must generate multiple variants for each kernel (e.g. by using C++ templates) for each reduction used in the calling Invoke or Iterate.

## 4. Compiling Orchestration

The orchestration constructs Invoke, Iterate and Pipe are used to invoke IrGL kernels. Since data-driven IrGL kernels often use worklists, the Iterate and Pipe also setup worklists. They also execute a series of IrGL kernels until the worklist is empty since iterative exection is a common pattern.

IrGL provides a default worklist object named `WL` that exposes `push`, `pop` and an iterator to each kernel. Therefore, `pop` and `push` are encoded in the AST as MethodInvocation on this object and do not appear as first-class AST nodes.

### 4.1 Worklist Mechanics

Kernels use worklists to manage work and as a means of communication of work between kernels. IrGL provides a default worklist to every kernel. A kernel may `push` values onto a worklist to enqueue work, and may `pop` values off the worklist to perform work, usually using a ForAll. IrGL worklists exhibit bulk-synchronous behaviour – work items pushed during an invocation cannot be popped in the same invocation.

```
1  Pipe Once {
2    Invoke identify_bad_triangles(mesh);
3    printf(``initial bad: %
4
5    // looping Pipe
6    Pipe {
7      Invoke refine(mesh);
8      ... // other mesh maintenance code
9
10     // only among newly created triangles
11     Invoke incremental_id_bad_triangles(mesh);
12   }
13
14   // sanity check
15   Invoke identify_bad_triangles(mesh);
16   printf(``final bad: %
17 }
```

**Listing 4.** Example of Pipe in DMR

```
1  Pipe Once {
2    Invoke A();
3
4    if(cond) {
5      Invoke B();
6    else {
7      Invoke C();
8    }
9  }
```

**Listing 5.** Example of Dynamic Piping

Worklists are created and managed by Iterate and Pipe constructs. Iterate is best illustrated by the BFS code in Listing 3. It creates a worklist, initially populated with $src$ and invokes the `BFS` kernel repeatedly until the worklist is depleted. After every invocation, code in `stmts` is executed. In this example, the `LEVEL` variable is incremented. The automatically created worklist is not available beyond the execution of the Iterate statement.

The Pipe statement establishes a shared worklist for the Iterate, Invoke and Pipe statements within it. A Pipe may execute once or loop until the worklist is empty. Inside a Pipe, the values pushed by an invocation of a kernel are forwarded to the next kernel in the pipe. Listing 4 illustrates the use of Pipe in the main loop of the DMR benchmark. After receiving an initial set of bad triangles, the inner Pipe iteratively refines the mesh, communicating the worklists between the two kernels inside the pipe.

The "flow" of worklists between kernels is not fixed at compile time. For example, Listing 5 is perfectly valid IrGL code. Depending on what `cond` evaluates to, the worklist produced by `A` may be consumed by either `B` or `C`.

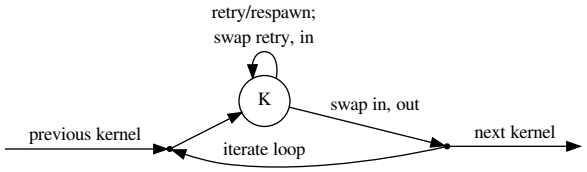


**Figure 1.** Flow control and worklist management for Iterate

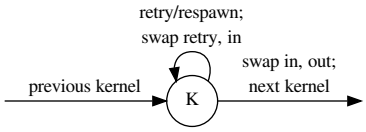


**Figure 2.** Flow control and worklist management for Invoke on a kernel that uses worklists

### 4.2 Compiling Pipe

The actual creation and communication of worklists between kernels is the responsibility of the Pipe/Iterate statements. Currently, the outermost Pipe (or Iterate) statement in a host Kernel creates a *pipe context*. All nested Pipe, Iterate and Invoke statements inherit this pipe context. In our implementation, the pipe context contains the incoming, outgoing and retry worklists named `in`, `out` and `retry` respectively. The `wlinit` attribute specifies the size of the worklist (`size`) and how the initial worklist is populated, which is implementation-dependent. For example, our compiler supports initializing worklists from a list of scalar expressions (WorklistInitializer) or from an array (WorklistInitializerFromArray).

When compiling an invocation to a kernel that reads, writes or iterates over the `WL` object, all `pops` are executed on the `in` worklist. Similarly all `pushes` are executed on the `out` worklist, using cooperative conversion where applicable to improve performance. Workitems to be retried are pushed into the `retry` worklist.

Figures 1 and 2 illustrate how control flows within a Iterate or Invoke. In general, flow is linear from the previous kernel to the next unless the kernel uses Iterate in which case the kernel is invoked repeatedly until no more items are left to process. Each invocation swaps the `in` and `out` worklists. If the kernel uses the `retry` worklist, it will be invoked repeatedly, but the `in` and `retry` worklists are swapped, while the `out` worklist remains the same.

Our compiler also sets up storage to store the return value when compiling Invoke and Iterate for kernels that use ReduceAndReturn.

Note that Iterate is syntactic sugar for a loop that wraps Invoke for kernels that do not use worklists. Iterate is essentially equivalent to a Pipe for kernels that do use worklists. Apart from terminating when the worklist is empty, it is possible to specify (in `extra_cond`) additional conditions that will cause the loop to exit even when the worklist is not empty. This extra condition may be combined with the empty worklist check using either And or Or.

### 4.3 Compiling Kernel Invocations

In the most general case, when the statement invoking the kernel, an Iterate or Invoke or Pipe, lies in a host kernel, a kernel invocation compiles down to a CUDA kernel launch. However, when the kernel invocation lies in the *control* kernel of an outlined Pipe, which is a CUDA `__global__`, it is compiled to a device kernel function invocation. Our compiler also supports the use of CUDA Dynamic Parallelism when launching kernels from control kernels, but the performance is poor, and it is not recommended.

Since IrGL kernels have no notion of threads, our compiler must also choose appropriate grid and thread block sizes for the CUDA launch. If SyncRunningThreads or Exclusive are not used in the kernel, then any grid size can be used, with our compiler using a fixed grid size calculated from the number of multiprocessors in the GPU. The use of these constructs requires that the grid size be chosen carefully as described earlier in Section 3.4. Without optimizations enabled, IrGL kernels also naturally compile down to *elastic* kernels [6] and so can run with any thread block size. However, when optimizations are enabled, the thread block sizes for a kernel may be constrained as we describe below.

Our compiler allows programmers to use the CUDA `__launchbounds__(maxthreadsperblock, minblocks)` annotation on individual kernels. The optional `minblocks` parameter is advisory and requests the compiler to achieve a residency of at least `minblocks` on each multiprocessor of the GPU. It is ignored by our compiler. The `maxthreadsperblock` parameter, on the other hand, informs the CUDA compiler that the kernel will not be launched with more than `maxthreadsperblock` which changes the behaviour of the register allocator. Attempting to launch a kernel with more than `maxthreadsperblock` will result in failure. Thus, `__launchbounds__` establishes an upper bound on the thread block size that can be used by our compiler.

Using nested parallelism or cooperative conversion also imposes a constraint on the thread block size that can be selected for a kernel. Essentially, both these optimizations make use of CUDA shared memory for communication with the size of shared memory used depending on the thread block size. Similarly, some libraries that we use internally use C++ template parameters to specialize for a statically specified thread block size. Such kernels are therefore limited to a fixed thread block size.

To summarize, IrGL kernels can fall into three categories depending on the thread block size they support. First are the *ElasticBlock* kernels, which can execute with any thread block size. Second are the *ShrinkableBlock* kernels, which place an upper bound on their thread block size. Finally, in the third category are the *FixedBlock* kernels, which can only execute with a fixed thread block size.

If iteration outlining is not used, the constraint on one kernel does not affect another kernel. However, when a Pipe

containing different kernels is outlined to the GPU and dynamic parallelism is *not* used, the thread block size chosen for the control kernel must satisfy the constraints on all kernels in the Pipe.

Since the maximum thread block size is limited by CUDA to 1024 on all GPUs we support, the set of possible thread block sizes for a kernel $k$, denoted by $T_k$, is finite. If $K$ is the set of kernels in a Pipe, then $T_{control}$ is simply:

$$T_{control} = \bigcap_{k \in K} T_k$$

Thus, the thread block size of the control kernel is simply the intersection of the domain sets for the constraint variables of each kernel. If this intersection is empty, then iteration outlining cannot be performed on this Pipe. If this intersection contains multiple values, our compiler chooses the highest value.

It is possible for ElasticBlock and ShrinkableBlock kernels to support different thread block sizes in different Pipes. However, for simplicity, our compiler picks a single thread block size for each kernel that is used at every invocation.

## 5. Conclusion

In this report, we have described the AST for IrGL and how it is lowered to CUDA. Our scope has been limited to the primary IrGL constructs since our intent was to provide a high-level overview of the process. We hope that this document will also be helpful to understand the organization of the IrGL compiler source that will be released separately. Among this document's omissions, we note absence of a discussion regarding the annotations supported by our compiler as well as its support for selecting optimizations at the Block level that allows a richer search space for auto-tuning, since these constructs are currently in flux.

## References


[1] J. Alglave, M. Batty, A. F. Donaldson, G. Gopalakrishnan, J. Ketema, D. Poetzl, T. Sorensen, and J. Wickerson. GPU concurrency: Weak behaviours and programming assumptions. In Ö. Özturk, K. Ebcioglu, and S. Dwarkadas, editors, *Proceedings of the Twentieth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '15, Istanbul, Turkey, March 14-18, 2015*, pages 577–591. ACM, 2015. ISBN 978-1-4503-2835-7. . URL `http://doi.acm.org/10.1145/2694344.2694391`.

[2] W. W. L. Fung, I. Singh, A. Brownsword, and T. M. Aamodt. Kilo TM: hardware transactional memory for GPU architectures. *IEEE Micro*, 32(3), 2012. .

[3] D. Merrill, M. Garland, and A. S. Grimshaw. Scalable GPU graph traversal. In *PPOPP 2012*. ACM, 2012. .

[4] R. Nasre, M. Burtscher, and K. Pingali. Morph algorithms on gpus. In *PPoPP '13*, PPoPP '13, 2013.

[5] L. Nyland and S. Jones. Understanding and using atomic memory operations. GTC 2013, 2013.

[6] S. Pai, M. J. Thazhuthaveetil, and R. Govindarajan. Improving GPGPU concurrency with elastic kernels. In V. Sarkar and R. Bodík, editors, *Architectural Support for Programming Languages and Operating Systems, ASPLOS '13, Houston, TX, USA - March 16 - 20, 2013*, pages 407–418. ACM, 2013. ISBN 978-1-4503-1870-9. . URL `http://doi.acm.org/10.1145/2451116.2451160`.

[7] A. Ramamurthy. Towards scalar synchronization in SIMT architectures. Master's thesis, The University of British Columbia, 2011.

[8] T. Sorensen and A. F. Donaldson. The hitchhiker's guide to cross-platform opencl application development. In *Proceedings of the 4th International Workshop on OpenCL, IWOCL 2016, Vienna, Austria, April 19-21, 2016*, pages 2:1–2:12. ACM, 2016. ISBN 978-1-4503-4338-1. . URL `http://doi.acm.org/10.1145/2909437.2909440`.

[9] V. Vineet, P. Harish, S. Patidar, and P. J. Narayanan. Fast minimum spanning tree for large graphs on the GPU. In *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS Conference on High Performance Graphics 2009*. ACM, 2009. .

[10] D. C. Wang, A. W. Appel, J. L. Korn, and C. S. Serra. The zephyr abstract syntax description language. In C. Ramming, editor, *Proceedings of the Conference on Domain-Specific Languages, DSL'97, Santa Barbara, California, USA, October 15-17, 1997*, pages 213–228. USENIX, 1997. URL `http://www.usenix.org/publications/library/proceedings/dsl97/wang.html`.

[11] S. Xiao and W. Feng. Inter-block GPU communication via fast barrier synchronization. IPDPS 2010. IEEE, 2010. .

```
module = Module(mod_stmts stmts)
mod_stmts = Kernel(identifier name, param *params,
                   bstmts stmts, str ret_type,
                   bool host, bool device)
          | CDeclGlobal(cdecl *decls, bool parse, bool dont_move)
          | CBlock(str *stmts, bool parse)
          | NOP()
          | Names(str *names)

param = (str type, identifier name)
cdecl = (str type, identifier name, str initializer)

bstmts = Block(stmts *stmts)

stmts = Block(stmts *stmts)
      | Assign(str lhs, str rhs)
      | For(identifier ndxvar, object iterator, bstmts *stmts)
      | ForAll(identifier ndxvar, object iterator, bstmts *stmts)
      | CFor(str init, str cond, str update, bstmts *stmts)

      | While(str cond, bstmts *stmts)
      | DoWhile(str cond, bstmts *stmts)

      | Atomic(str lock, str lockndx, stmts *stmts, bstmts *fail_stmts)
      | Exclusive(object_list objs_lock, stmts *stmts, bstmts *fail_stmts)

      | Retry(str args, bool merge)
      | Respawn(str args)
      | CBlock(str *stmts, bool parse)
      | CDecl(cdecl *decls, bool parse, bool dont_move)

      | If(str cond, bstmts *true_stmts, bstmts *false_stmts)

      | Invoke(aggr_func aggr, identifier kernel, str *args)
      | Iterate(term_cond cond, aggr_func aggr, identifier kernel,
                str *args, worklist_initializer wlinit, bstmts *smts,
                extra_cond extra_cond)
      | ReduceAndReturn(str value)

      | SyncRunningThreads()
      | LocalBarrier()

      | Pipe(bstmts *stmts, worklist_initializer wlinit, bool once)

      | Expr(expr e)

expr = CExpr(str expr, bool parse)
     | MethodInvocation(str obj, identifier method, str obj_type, str *args)

object_list = Tuple(identifier object, item_list *items)

item_list = Array(int size, identifier name)
          | ArrayIterator(identifier array, str start, str end, str step)

term_cond  = While | Until
aggr_func  = Any | All
extra_cond = And(str cond) | Or(str cond)

worklist_initializer = WorklistInitializer(str size, str *initial)
                     | WorklistInitializerFromArray(str size, identifier array, str array_size)
```

**Listing 6.** IrGL Abstract Syntax Tree in ASDL